arXiv:1401.2649v1 [stat.AP] 12 Jan 2014

Personalised Medicine: Taking a New Look at the Patient

Marco Scutari

Genetics Institute, University College London, United Kingdom.

Personalised medicine strives to identify the right treatment for the right patient at the right time, integrating different types of biological and environmental information. Such information come from a variety of sources: omics data (genomic, proteomic, metabolomic, etc.), live molecular diagnostics, and other established diagnostics routinely used by medical doctors [6]. Integrating these different kinds of data, which are all high-dimensional, presents significant challenges in knowledge representation and subsequent reasoning [4, 18]. The ultimate goal of such a modelling effort is to elucidate the flow of information that links genes, protein signalling and other physiological responses to external stimuli such as environmental conditions or the progress of a disease.

Omics data, which include single-nucleotide polymorphisms (SNPs), protein and gene regulatory networks, are investigated using high-throughput platforms such as the ones developed for the Human Genome project [2, 14]. Systems biology studies the relationships among the elements in omics data as they change in the presence of genetic and environmental perturbations, extending techniques that were previously used on a smaller scale [13]. Such knowledge can improve our ability to understand and predict the behaviour of complex biological systems, but requires careful handling in integrating different sources. On its own, each type of data often contains too much noise for single biological signals to be identifiable, much less their interplay. Pooling the information available across omics data (e.g. sequencing and expression information about relevant genes, possibly under different treatment regimens) provides an option to increase statistical power and produce reliable knowledge representation models [10].

The role of live molecular diagnostics, and to some extent of traditional diagnostics, is to complement omics data with longitudinal measures of the patient's condition that are easier and cheaper to collect. Several examples of the modelling and implementation techniques involved are covered in the previous sections. Integrating such diagnostics is essential because genetic information correlate only imperfectly with protein levels [8], which in turn are very noisy predictors of most pathologies.

Applications of personalised medicine fall roughly in three groups. Firstly, drug discovery and development can be made more efficient and effective [6]. On the one hand, omics data can provide feedback at early stages of drug discovery by replacing the traditional trial-and-error approach with a hypothesis-driven one based on a formal knowledge representation model. On the other hand, omics data can also be used to improve clinical trial design by guiding patients selection and stratification based on predicted drug toxicity and non-responders profiles. This is likely to prove more effective than defining populations in terms

of race or ethnicity, since only $5-10\%$ of the total human genetic variance occurs between different ethnic groups [1] and boundaries between different populations are often not clear.

Secondly, several aspects of the diagnostic process can be improved. For example, the normal behaviour of a biological system can be better defined at the molecular level than using non-specific clinical signs. As a result, pathologies can be classified with greater precision based on a molecular taxonomy [6]; previously unknown differences have been highlighted in breast cancer [7] and leukaemia [17] in this way. Furthermore, genetic tests need to be improved in their sensitivity and specificity; they are challenging to perform reliably and interpret correctly, and they focus predominantly on rare diseases [9].

Thirdly, personalised medicine allows treatment for many diseases to be tailored to each patient to an unprecedented degree. For example, adverse reactions to specific compounds can be predicted with greater accuracy, and non-responders can be identified without actually starting a therapy that may or may not be effective.

To investigate and implement personalised medicine in practice, many challenges need to be overcome at the modelling level; some of them will be covered in the following chapter. First and foremost, a working knowledge representation model must be established to facilitate reasoning on high-dimensional, heterogeneous data. Currently, probabilistic graphical models (Bayesian networks in particular) seem to be a popular approach [3, 5, 12]. Their ability to provide at the same time an intuitive understanding of the data to biologists and medical doctors (through the graph structure) and a rigorous probabilistic framework to statisticians and computer scientists makes them an ideal tool for this task.

Moreover, specific distributional assumptions are required to accurately describe both omics and diagnostics data effectively. Gaussian and discrete Bayesian networks from classic literature [16] present important limitations in modelling omics data, as do more general models such as chain graphs. For instance, assuming normality for gene expressions will almost certainly result in a biased model, because expression levels are usually highly skewed. Likewise, ignoring the ordering of the alleles in SNP data disregards information which is known to be fundamental in quantitative genetics. Ideally, probabilistic assumptions should also support the inclusion of available prior information from different sources, as in Schadt et al. [15].

A related issue is the computational complexity of both model estimation and subsequent inference, which poses severe limits to the use of flexible distributional assumptions in Bayesian networks and to the scope of the questions these networks can answer. The use of prior information can speed up model estimation by reducing the set of the models under consideration, even though it may introduce bias as well if the phenomenon we are modelling is not well understood. Another possible solution is to perform feature selection as a preprocessing step, thus speeding up inference as well. In the context of Bayesian networks, Markov blankets provide a natural way to do so while retaining as much information as possible [11]. However, given the complexity of the data

used in personalised medicine, the cost of feature selection is often as high as that of model estimation.

In conclusion, while there are many open problems to address, an effective use of knowledge representation is crucial in implementing reliable personalised medicine protocols. Omics and other established diagnostics provide a wealth of data, which calls for appropriate modelling spanning techniques from statistics, computer science and quantitative biology.

Bibliography

[1] Cavalli-Sforza, L.L., Feldman, M.W.: The Application of Molecular Genetic Approaches to the Study of Human Evolution. Nature Genetics 33, 266–275 (2003)

[2] Collins, F.S., Morgan, M., Patrinos, A.: The Human Genome Project: Lessons from Large-Scale Biology. Science 300(5167), 186–290 (2003)

[3] Cooper, G.F., Hennings-Yeomans, P., Visweswaran, S., Barmada, M.: An Efficient Bayesian Method for Predicting Clinical Outcomes from Genome-Wide Data. In: AMIA Annual Symposium Proceedings. pp. 127–131 (2010)

[4] Emmer-Streib, F.: Personalized Medicine: Has it Started Yet? A Reconstruction of the Early History. Frontiers in Genetics 3(313), 1–4 (2013)

[5] Friedman, N., Linial, M., Nachman, I.: Using Bayesian Networks to Analyze Expression Data. Journal of Computational Biology 7, 601–620 (2000)

[6] Ginsburg, G.S., McCarthy, J.J.: Personalized Medicine: Revolutionizing Drug Discovery and Patient Care. Trends in Biotechnlogy 19(12), 491–496 (2001)

[7] Golub, T.R., Slonim, D.K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J.P., Coller, H., Loh, M.L., Downing, J.R., Caligiuri, M.A., Bloomfield, C.D., Lander, E.S.: Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring. Science 286(5439), 531–537 (1999)

[8] Gygi, S.P., Rist, B., Gerber, S.A., Turecek, F., Gelb, M.H., Aebersold, R.: Quantitative Analysis of Complex Protein Mixtures Using Isotope-Coded Affinity Tags. Nature Biotechnology 17, 994–999 (1999)

[9] Hamburg, M.A., Collins, F.S.: The Path to Personalized Medicine. The New England Journal of Medicine 363, 301–304 (2010)

[10] Ideker, T., Thorsson, V., Ranish, J.A., Christmas, R., Buhler, J., Eng, J.K., Bumgarner, R., Goodlett, D.R., Aebersold, R., Hood, L.: Integrated Genomic and Proteomic Analyses of a Systematically Perturbed Metabolic Network. Science 292(5518), 929–934 (2001)

[11] Koller, D., Sahami, M.: Toward optimal feature selection. In: Proceedings of the 13th International Conference on Machine Learning (ICML). pp. 284–292 (1996)

[12] Mourad, R., Sinoquet, C., Leray, P.: A Hierarchical Bayesian Network Approach for Linkage Disequilibrium Modeling and Data-Dimensionality Reduction Prior to Genome-Wide Association Studies. BMC Bioinformatics 12(16), 1–20

[13] Sachs, K., Perez, O., Pe'er, D., Lauffenburger, D.A., Nolan, G.P.: Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data. Science 308(5721), 523–529 (2005)

[14] Sawicki, M.P., Samara, G., Hurwitz, M., Jr., E.P.: Human Genome Project. The American Journal of Surgery 165(2), 258–264 (1993)

[15] Schadt, E.E., Lamb, J., Yang, X., Zhu, J., Edwards, S., Guhathakurta, D., Sieberts, S.K., Monks, S., Reitman, M., Zhang, C., Lum, P.Y., Leonardson, A., Thieringer, R., Metzger, J.M., Yang, L., Castle, J., Zhu, H., Kash, S.F., Drake, T.A., Sachs, A., Lusis, A.J.: An Integrative Genomics Approach to Infer Causal Associations Between Gene Expression and Disease. Nature Genetics 37(7), 710–717 (2005)

[16] Scutari, M., Strimmer, K.: Introduction to Graphical Modelling. In: Balding, D.J., Stumpf, M., Girolami, M. (eds.) Handbook of Statistical Systems Biology. Wiley (2011)

[17] Waring, J.F., Ciurlionis, R., Jolly, R.A., Heindel, M., Ulrich, R.G.: Microarray Analysis of Hepatotoxins in Vitro Reveals a Correlation Between Gene Expression Profiles and Mechanisms of Toxicity. Toxicology Letters 120(1–3), 359–368 (2001)

[18] Weston, A.D., Hood, L.: Systems Biology, Proteomics, and the Future of Health Care: Toward Predictive, Preventative, and Personalized Medicine. Journal of Proteome Research 3(2), 179–196 (2004)